Intel® Storage System SSR212MA

Software Release Notes

Revision 1.5

May, 2007

Storage Systems Technical Marketing

Revision History

Date	Revision Number	Modifications
December, 2005	1.0	1st Release copy.
January, 2006	1.1	Added 6 additional issues: 2.1, 3.8, 3.9, 3.10, 11.12, 12.4.
June, 2006	1.2	Added 6.3 SP1 updates & issues, noted by "(SP1)" in the issue title.
August, 2006	1.3	Added 6.5 updates & issues
January, 2007	1.4	Added 6.6 updates & issues
May, 2007	1.5	Added 6.6 SP1 updates & issues

Disclaimers

THE INFORMATION IN THIS DOCUMENT IS PROVIDED "AS IS" WITH NO WARRANTIES WHATSOEVER, INCLUDING ANY WARRANTY OF MERCHANTABILITY, FITNESS FOR ANY PARTICULAR PURPOSE, OR ANY WARRANTY OTHERWISE ARISING OUT OF ANY PROPOSAL, SPECIFICATION, OR SAMPLE.

Information in this document is provided in connection with Intel® products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, life saving, or life sustaining applications.

Intel retains the right to make changes to its test specifications at any time, without notice.

The hardware vendor remains solely responsible for the design, sale and functionality of its product, including any liability arising from product infringement or product warranty.

Table of Contents

1 Introduction

The following Release Notes provide information about current limitations in this 6.6 SP1 release of the Storage System SAN software, and 6.6 SP1 release of the Storage System Console (SSC) software for SSR212MA.

2 Summary of Issues Fixed in SAN Software 6.6 SP1

Enhancement – Improved warning dialogs when deleting volume(s) and/or snapshot(s)
Rebooting a storage module with Microsoft Cluster nodes connected causes NTFS Delayed Write failure and/or ftdisk Windows event log error messages.
Storage server process may restart while under heavy load and another storage module is restarted
Unable to enable flow-control on the storage module
The storage server takes longer to start up than previous Storage System Software releases
Microsoft Cluster Resources may go offline when using the Storage System Software DSM for MPIO
Refreshing the Hardware Information report continually for over 200 times can cause process failure
Reduce the number of clock interrupts so that more CPU cycles are available for I/O
DSM services does not allow logging off one iSCSI session per target on the Microsoft cluster node that is not hosting any resources
Add Microsoft iSCSI 2.03 to the Virtual IP Load Balancing compliance list
The storage server can sometimes use all the physical memory on a high capacity storage module, causing memory exhaustion
Reworked Storage System Software predictive drive failure logic to adhere to the SMART standards of the drive vendor
Enhanced iSCSI reserve/release logic to perform operations faster (VMware Patch 1000500)
VMware robustness improvements added
Fixed slow memory leak during Storage system Console operations that query volume availability
Adjustments to platforms running the Linux 2.6 kernel that can result in a minor sequential write performance gain
Reduced manager response time for new database operations to prevent potential volume availability issues
Fixed manager compatibility issue when running a mixed management group (version 6.5 and 6.6)
Added logic to assure that hotfix binaries are removed when applying a service pack

3 Upgrading to 6.6 SP1

Important Notice!

For all SSR212MAs currently running 6.6.00.xxxx, Patch 10007-01 (reboot not required) must be installed before upgrading to version 6.6 SP1. Please note that Patch 10007-02 is already included in 6.6 SP1 software release so no installation is required.

Platforms Supported for This Release

Upgrades to Release 6.6 SP1 are available for SSR212MA. Please use the following procedures to upgrade storage modules from 6.5.xx.xxxx or 6.6.xx.xxxx to 6.6 SP1.

If you see the following type of message during the upgrade process, please call customer support.

```
UPGRDE: 6.6.xx.xxxx.main—Upgrade will now be aborted.

Thu May 19 22:34:56 GMT 2005: UPGRADE: 6.6.xx.xxxx.main aborted
```

3.1 Special Feature Key Upgrade Procedure for This Release

If you are using a feature key for add-on features and applications in releases prior to Release 6.5, these features keys must be reapplied to individual SSMs when upgrading to Releases 6.5 or above. This is due to a different feature key construct and syntax in the newer releases. If a reapplication is necessary you will be prompted near the end of the upgrade procedure, as described in the upgrade procedures below.

3.1.1 Feature Key Overview

Storage System Software 6.5 and above

With Storage System Software 6.5 and above, keyed features are enabled per SSM. Prior to Release 6.5, keyed features are enabled per management group. Feature keys are required for the following add-on features and applications:

- Scalability Pak
- Configurable Snapshot Pak
- Remote Data Protection Pak

Customers may use these add-on features and applications without a feature key, but are limited to a 30-day trial period. After the 30-day trial period, if a feature key is not purchased, any volumes and snapshots associated with add-on features or applications will become unavailable until a feature key is purchased and applied.

3.1.2 Prerequisites

- If you are running a software version earlier than 6.5.00, you must first upgrade to 6.5. You can then upgrade to Release 6.6 SP1 directly.
- If any iSCSI volumes are in use, stop any activity to those volumes and unmount or log off before beginning the upgrade. The storage module reboots as part of the upgrade process. Consequently, your volumes may go off-line, depending on your configuration.
- If you are running iSCSI load balancing, you must have Virtual IP Addresses configured.
- Ensure that you are using version 6.6.01.4006 of the Console before upgrading the Storage System Software on the storage modules.

3.1.3 Download The Upgrade Components

1. Download the upgrade components to a temporary location.
2. Download the latest Storage System Console (6.6.01.4006) from the SSR212MA support site or from your IBL account.
3. Download the appropriate Storage System Software upgrade package(s) from the SSR212 MA support site or from your IBL account.

Platform	For Storage System Software Version	Use Upgrade File
SSR212MA	6.5 or greater	6.5.x or 6.6.x to 6.6.01.6007.20070326.SSR212MA.upgrade

3.2 Install the Storage System Software 6.6 SP1 Console

1. Install the Storage System Software 6.6.01.4006 Console and discover storage module(s) on the network.
2. Use the Console to install the Storage System Software 6.6.SP1 upgrade. If you do not have a direct path to the 6.6 SP1 release, you may be required to first upgrade the software on the SSM(s) to a version that can then upgrade to 6.6 SP1

3.2.1 Best Practice

- Virtual IP Addresses - If a Virtual IP (VIP) address is assigned to a storage module in a cluster, the VIP storage module needs to be upgraded last. The VIP storage module is shown in a field in the clusters detail tab.
  - First upgrade the non-VIP storage modules that are running managers one at a time.
  - Then upgrade the non-VIP non-manager storage modules.
  - Lastly, upgrade the VIP storage module.

- Remote Copy – If you are upgrading management groups with Remote Copy associations, you must upgrade the remote management group first. If you upgrade the primary group first, Remote Copy will stop working temporarily, until both the primary management group and the remote group have finished upgrading.

3.2.2 Selecting The Type of Upgrade

The Storage System Software Console supports two types of upgrades, as shown in the figure below.

- One-at-a-time (recommended) - this is the default and only method if the storage modules exist in a management group.
- Simultaneous (advanced) - this allows you to upgrade multiple storage modules at the same time if they are not in a management group. Use this for new storage modules and/or re-configured storage modules.

1. Select from the list which storage modules to upgrade.
2. Select the type of upgrade.
3. Click Install.

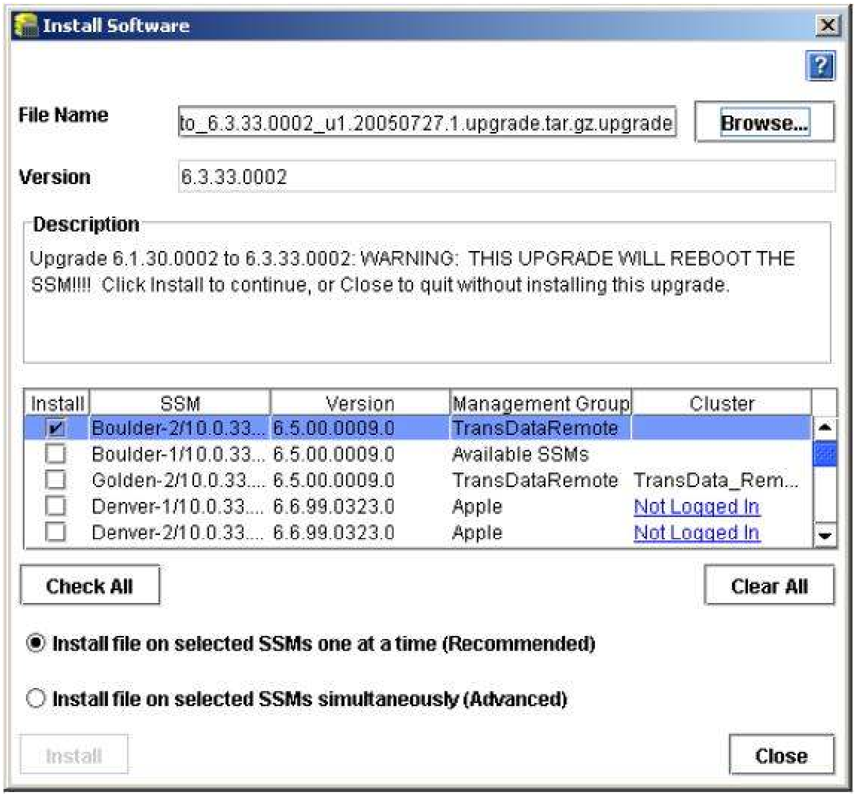

File Name: to_6.3.33.0002_u1.20050727.1.upgrade.tar.gz.upgrade

Version: 6.3.33.0002

Description

Upgrade 6.1.30.0002 to 6.3.33.0002: WARNING: THIS UPGRADE WILL REBOOT THE SSM!!!! Click Install to continue, or Close to quit without installing this upgrade.

Install	SSM	Version	Management Group	Cluster
☑	Boulder-2/10.0.33...	6.5.00.0009.0	TransDataRemote	
□	Boulder-1/10.0.33...	6.5.00.0009.0	Available SSMs	
□	Golden-2/10.0.33....	6.5.00.0009.0	TransDataRemote	TransData_Rem...
□	Denver-1/10.0.33....	6.6.99.0323.0	Apple	Not Logged In
□	Denver-2/10.0.33....	6.6.99.0323.0	Apple	Not Logged In

● Install file on selected SSMs one at a time (Recommended)

○ Install file on selected SSMs simultaneously (Advanced)

After The Storage Module Reboots

After each storage module upgrade, during the management group health check, you may see messages such as “Waiting for MG1 to come up. The issue is –An SSM is down.”

The storage module is not down. It is actually resyncing with the other storage modules in the management group.

3.2.3 Verify Management Group Version

You must verify the management group version only when upgrading 6.5 to 6.6 or 6.6 SP1. Management group will remain at 6.6 if upgrading from 6.6 to 6.6 SP1. To verify the management version please take the following steps.

1. Select the management group in the navigation window
2. Click the Register tab to view the management group version

After the upgrades are complete, the Console attempts to upgrade the management group version. The rules for this operation are as follows

- Upgrade management group version to 6.6 if all storage module serial numbers (eth0 MAC address) are in the list of known serial numbers included in the upgrade file
- Do not upgrade the management group version to 6.6 if any of the storage module serial numbers are unknown.
- If you get the following message (or similar message), the management group version upgrade has not completed. Until the management group is upgraded to 6.6.

If you get the following message (or similar message), the management group version upgrade has not completed. Until the management group is upgraded to 6.6, you will not be able to take advantage of all the 6.6 features.

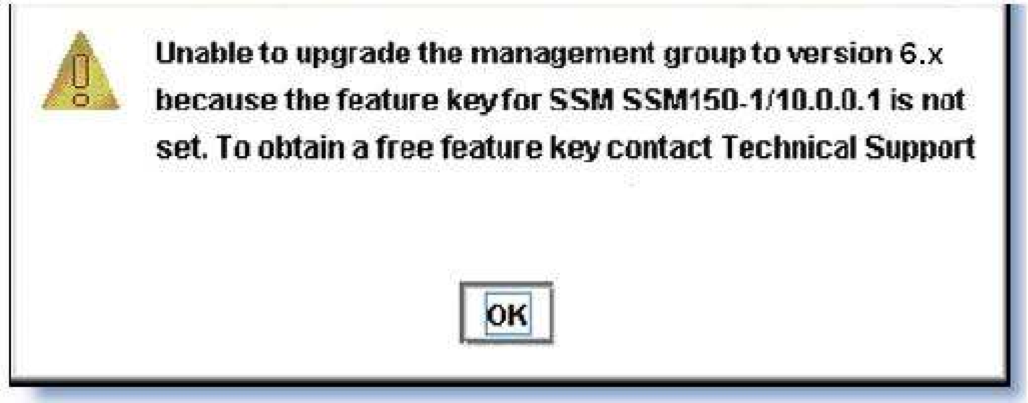

Unable to upgrade the management group to version 6.x because the feature key for SSM SSM150-1/10.0.0.1 is not set. To obtain a free feature key contact Technical Support

4 Current SAN & SSC Software Limitations

4.1 Storage System Console (SSC)

4.1.1 Storage System Console Fails to Install On linux for 6.6.xx.xxxx Release of The SSC.

Issue

When downloading the installer for the Console from the vendor's FTP site, the FTP program reports that the download completed succesfully. However, when you run the installer, you receive an error message indicating that a Java error occurred and the installtion cannot continue.

This occurs because some FTP programs may not download the complete installation package. You can verify that the download was complete by comparing the MD5 checksum of the file that was downloaded with the MD5 checksum that is published on the FTP site.

Fix

Upgrade the FTP client you are using or use a different FTP client.

4.2 Upgrades

4.2.1 Upgrade Post-Qualification May Grab Focus Every 20 Seconds

Issue

During a software upgrade, the Storage System Console may come to the front of other windows open on the desktop and may grab focus as well.

Workaround

None.

4.2.2 Upgrading Storage Modules and Management Groups May Take Some Time

Issue

Upgrading a storage module from 6.5.xx to 6.6.xx takes from 30-45 minutes depending upon the specific platform and configuration. Additionally, after the storage modules are upgraded, they have rebooted, and have all been found on the network in the Console, the management group health check may take up to another 10 minutes. During the management group health check you may see messages such as “Waiting for MG1 to come up. The issue is – An SSM is down.” The storage module is not down. It is actually resyncing with the other storage modules in the management group.

Workaround

None.

4.2.3 On SUSE Linux Enterprise 10, Storage System Console Installation Stops with Error

Issue

You cannot install or run Storage System Console in SLES 10.

Workaround

1. Edit the bin file (CMC_6.6.00.0099_installer_Linux.bin)
2. Comment out the following line

   Export LD_ASSUME_KERNEL = 2.2.5

   As

   #xport LD_ASSUME_KERNEL = 2.2.5
3. Now the installation completes successfully
4. After the installation completes, edit the following file as above.

   /opt/LeftHandNetworks/UI/LeftHand_Networks_Centralized_Management_Console

5. Now the Console will start and all the applications will run normally.

4.2.4 Upgrading Storage System Software Fails With Error Code 169

Please call customer support for help with the upgrade.

4.3 Storage System Module

4.3.1 Windows Firewall Prevents Storage Module Discovery In The Storage System Console

After upgrading the Storage System Software to version 6.6 or higher from version 6.5 or lower, the Storage System Console fails to discover storage modules.

Workaround

1. Determine if Windows Firewall is running.
2. If Windows Firewall is running, disable it.

4.3.2 How to Correctly Identify a Faulty Power Supply

The Intel® Storage System SR212MA ships with only one power supply; therefore, only one power supply is listed in the passive report. Status will be either normal or faulty.

Issue

In a system that has been upgraded to add a redundant module, if a power supply is not working properly, the storage console passive report it will report power supply status as faulty. The faulty module number will not be identified.

Workaround

To identify the storage module with a faulty power supply.

1. On the Module Information tab, click Set ID LED On.
2. The ID LED on the left front of the module illuminates a bright blue. Another ID LED is located on the back of the module on the right side under the empty slot.
3. Go to the back of the storage module and look at the two power supplies.
4. A green LED will be illuminated on the working power supply and an amber LED on the faulty power supply.
5. Replace the faulty power supply.

Note: To ensure redundancy, the two power cords must be connected to separate and independent power sources.

4.3.3 Rebooting the Storage Module While RAID Is Rebuilding Causes Reboot to Take Up To 20 Minutes

Issue

If you reboot the storage module while RAID is rebuilding, the reboot can take up to 20 minutes to complete.

Cause

The lower the priority setting of the RAID, the longer it will take the reboot to complete.

4.3.4 Repair Storage Module Stalls When Attempting To Remove the Storage Module

Workaround

1. Close the Storage System Console and reopen it. The storage module has moved from the cluster to the management group and the ghost storage module is in the cluster.
2. Remove the storage module from the management group.

Cause

The lower the priority setting of the RAID, the longer it will take the reboot to complete.

4.4 RAID and Disk Management

4.4.1 Why RAID May Go Off If a Foreign Drive Is Inserted Prior To Powering Up the SSR212MA

Issue

If the storage module powers up with a drive that does not belong to the RAID configuration, data corruption may occur causing RAID to go off and preventing the storage module from coming online. Replacing the original drive may not result in RAID going to normal. Data may be lost on this storage module in this case.

Workaround

Drive replacement should ALWAYS be done using the Console. Select the drive to replace, click power-off, insert a new drive, click power-on, and then click add-to-RAID.

4.4.2 Swapping One or More Disks across Controllers Causes Data Loss

If the storage module powers up with one or more drives foreign to the configuration of a controller, data corruption occurs.

Issue

The storage module is moved a different physical location. Before the move, the storage module is powered down and all drives are removed. While replacing the drives back in the drive bays, one or more drives are accidentally inserted into slots handled by a different controller. When the storage module is powered up, data corruption occurs.

Workaround

Labels should be added to drive carriers when first installed. If this has not been done, label the drives before removing them so that you can replace them in the correct bays.

4.4.3 What to Do When A Cache Corruption Alert Is Received

Issue

Cache corruption can occur if the storage module is powered down while there is data in the RAID cache. If the storage module stays powered-off long enough (more than 72 hours), data in the cache will be corrupted. When the storage module powers back up, the cache corruption is detected and an alert is posted indicating the cache is corrupt. The storage module will not be allowed to come online in order to prevent corruption within the cluster. A “storage module down” alert will also be posted. Please note that data on the storage module had been lost in the case and must be rebuilt from the cluster assuming replication was configured.

Workaround

To resolve this issue, please contact support

4.4.4 Rebuilding RAID 5 Takes To Long When Minimum Setting is 1

Issue

The default setting for the minimum RAID rebuild rate is 1. This setting may cause RAID 5 rebuild to take too long.

Workaround

Increase the minimum rebuild rate to a value of 10 or greater. The following guidelines describe the effects of the RAID rebuild rates.

- Setting the rate high is good for rebuilding RAID quickly and protecting data; however, it will slow down user access.
- Setting the rate low maintains user access to data during the rebuild.

4.4.5 When Replacing a Disk, if New Disk is seated improperly Disk Status Displays DMA Off With Yellow Exclamation Icon

Issue

A disk is replaced in an SSR212MA. After the RAID rebuild is complete, the disk status displays DMA Off. This status occurs due to an improperly seated disk.

Workaround

Repeat the procedures for replacing the disk, paying careful attention to reseat the disk properly in the drive bay. After the RAID rebuild is finished, the disk status should be correct.

4.4.6 Removing Drive from SSR212MA without First Removing Disk from RAID Requires Rebooting the SSR212MA to Recover from Degraded Mode

Issue

If a drive is removed without first removing it from RAID in the Console, RAID becomes degraded and the SSR212MA becomes inaccessible.

Workaround

1. Re-insert the drive.
2. Reboot the module.
3. Add the disk to RAID. RAID will start rebuilding after the drive is powered on.

4.4.7 SSR212MA Becomes Inaccessible after RAID BIOS Deletion

Issue

After deleting RAID configuration from the RAID BIOS in the MA system, the system can no longer be accessed, both through SSC connection or through the serial link. The system will boot up normally and responde to pings from the network, however, cannot be access through the SSC software. When trying to connect through the serial link you will get error: “Error opening a connection”.

Workaround

There is currently no workaround for this issue.

4.4.8 No Warning If Remove and Re-Add Disk to RAID 0.

Issue

SSR212MA is configured with RAID 0. While the SSR212MA is running, user manually removes any disk from the SSR212MA. On the Disk Setup window the disk status is “Off or missing.” On the RAID Setup window, RAID status is Normal.

This Issue occurs when the disk is removed while there is no activity to the volume. As soon as any activity to that volume occurs, such as a client attempting to read or write data, then the volume becomes unavailable.

4.4.9 Single Drive Error

Issue

A drive may become unavailable (especially with Hitachi 250GB drive), causing the RAID status to go Degraded or Off, depending on the RAID configuration.

Workarounds

The following three options should be tried in order. If one does not fix the problem, try the next one.

- Reseat the drive, using the instructions in the User manual or Online Help. If the drive does not start rebuilding and the drive status shows Inactive in the Disk Setup tab, select the drive and click Add to RAID.
- Reboot the storage module. The drive comes online and begins rebuilding.
- Replace the drive and rebuild the array.

4.4.10 Storage Module Becomes Unavailable After Replacing A Disk in RAID 0

Issue

After replacing a drive in a storage module configured in RAID 0, the storage module becomes unavailable.

Workarounds

Reboot the storage module.

4.5 Network Management

4.5.1 Storage Traffic Is Using The Wrong Network Interface Card (NIC)

Issue

You may see Storage System Software traffic on NICs other than the designated one. This is unavoidable when two or more NICs are assigned IP addresses in the same subnet. It can occur in any configuration where hosts are configured with multiple NICs.

Workaround

Assign “public” adapters, intended for servicing users, to a subnet distinct from storage adapters.

4.5.2 Configuring the SAN on A Private vs. Public Network

Issue

The recommended best practice is to isolate the SAN, including Console traffic, on a separate network. If the SAN must run on a public network, use a VPN to secure data and Console traffic.

Workaround

None.

4.5.3 When Jumbo Frames Are Used With Incorrect Network Setup, Management Group Is Rendered Unusable

Issue

If a management group is created with a storage module that has an improper network configuration, such as an active-backup bond with one network interface connected to a 10/100 switch and another connected to a GigE switch, the management group losses connectivity.

Also, if such a bond is deleted, the storage module permanently loses connectivity.

Workaround

A proper network configuration is required when creating a bond; that is, both network interfaces must be connected to appropriate switches. For more information, see the *Intel® Storage System Software User Manua*l, “Configuring NIC Bonding.”

4.5.4 NIC Cards Generate DHCP Requests and Incorrect MAC Addresses

Issue

NIC cards may send unwanted DHCP requests which appear to come from a NIC with a MAC address that does not match either NIC.

Workaround

The DCHP requests are coming from an IPMI network port that is currently not being used. These DHCP requests are not harmful, however, to disable this port please contact customer support.

4.6 Reporting and SNMP

4.6.1 "NVRAM Card = Corrupt" Alert Generated When the Storage Module is Restarted After Being Shut Down for Some Time

Workaround

To resolve this issue, please contact support.

4.6.2 NTP and SNMP Setting Are Not Retained After DOM Replacement

Workaround

Reconfigure the NTP and SNMP settings in the Storage System Console.

4.6.3 Hardware > Log Files Tab Presents an Un-refreshed List OF Log Files

Issue

If you select > Storage Module Tasks > Edit Configuration > Hardware > Diagnostics > Run Tests. The file /var/log/hpadu.log is created, but may not show in the Log files tab window of Edit Configuration. This is because the window is incorrectly refreshed, not because the log file is absent.

Workaround

Log out of the storage module and log back in to refresh the file list.

4.7 Management Groups

4.7.1 Restoring a Management Group Configuration Fails Because the Configuration Backup Does Not Reflect AN IP Address Change Made To AN SSR212MA

Issue

This Issue can occur when the IP address change is applied to an SSR212MA, but does not get stored in the system before the management group configuration is backed up.

Workaround

If you do restore a management group configuration with an incorrect IP address, take the following steps:

1. Log in to the SSR212MA with the incorrect IP address and change it back to the IP address stored in the configuration backup file.
2. Complete the management group restoration.
3. When the management group is successfully restored, change the IP address of the SSR212MA to the desired address.

Best Practice to Prevent This Problem

1. In the Edit Configuration window, change the IP address of the SSR212MA.
2. In the Console Network View, select the management group.
3. Right-click and select Backup Configuration of Management Group.
4. In the Back up Configuration window that opens, scroll down to the section labeled SSR212MAs.
5. In the SSR212MAs section, verify that the Communication IP is the new IP address. If it is not the new IP address, then click OK to cancel out of the Backup Configuration window.

Wait for a few minutes and then repeat steps 3 through 5. When the correct new IP address appears, select Back up Configuration. The management group configuration is backed up with the correct IP address.

4.7.2 Management Group Restoration Fails When Storage Module Have Different Software Versions (6.5 and 6.6) And Are Running Managers

Solution

1. Upgrade all the storage modules to the same version of the software.
2. Backup the management group again, after all storage modules have been upgraded. Then future management group restorations will work.

Or

1. On storage modules running version 6.5, you can start and stop managers (as long as quorum can be maintained) so that only the 6.5 storage modules in the management group are managers.
2. Backup the management group configuration again and restore.

The key to implementing either solution is to take a backup of the management group configuration again before attempting to restore.

4.7.3 Management Group IP Addresses Are Misconfigured Which Causes Manager communication Problems When A Storage Module Reboots

Issue

Creating a management group using the wizard results in an incorrectly configured unicast list. When a storage module in the management group reboots, the management group may seem to have lost quorum.

Workaround

For each storage module in the management group:

1. Open the Edit Configuration window.
2. Select the TCP/IP configuration category and select the Communication tab.
3. Click Update (at the bottom right corner)
4. Close the Edit Configuration window.

4.7.4 After Deleting A Virtual Manager, Removing A Storage Module From the Management Group Fails The First Time

Issue

You attempt to remove a storage module from a management group. The storage module is running a manager. The management group has a Virtual Manager added and started. First you stop the Virtual Manager. Then you delete the Virtual Manager from the management group. Next you stop the manager on the storage module. Now, when you remove the storage module from the management group, the process seems to complete, but the storage module does not actually get removed.

Workaround

Perform the remove operation again on the storage module. The straoge module should be removed from the management group.

4.7.5 License Evaluation message Opens When Using Mangement Group Wizard, Even Though All Storage Modules Are Licensed.

Workaround

Click through the message and ignore it.

4.7.6 Storage System Console Falsely Indicates Loss Of Quorum For Management Group

Issue

If the storage module running the manager through which the Console is logged in, loses contact with the rest of the managers, the Console reports loss of quorum and the entire management group flashes red. However, the issue may be solely with the specific manager, and not the entire management group, which may still have quorum.

Workaround

1. Review the management group Details tab to find storage modules running managers that display the manager status as normal. Note the IP addresses of those storage modules.

2. From the menu bar, select Find > Clear All Found Items. All the management groups and available storage modules are cleared from the navigation window.
3. Now use the Find Modules wizard or the Find menu to search for one of the IP addresses you wrote down in step 1.
4. Log in to that storage module and review the quorum status and the manager status. If it reports no quorum, try the next manager IP address you wrote down.

   By checking the “normal” managers, you may find that it is a single manager having problem, and not the entire management group. You may find that the management group is operating and volumes are still on line.

   If this is the case, you can proceed to troubleshoot the manager issue.

4.8 Clusters

4.8.1 If Incorrect Virtual IP Information Is Entered, SSR212MAs Go Offline and Volumes Become Unavailable

Issue

When configuring VIP for a cluster, entering incorrect information for any of the components (IP Address, Subnet Mask and Default Gateway) causes the SSR212MAs in that cluster to go down and any volumes associated with the cluster to become unavailable.

Workaround

The iSCSI VIP must be in the same subnet as all the SSR212MAs in the cluster.

1. Enter the correct information for the Virtual IP configuration.
2. Reboot the SSR212MAs in the cluster.

.

4.9 Volumes & Snapshots

4.9.1 Snapshot Schedules Do Not Adjust For Daylight Savings Time

Issue

When snapshot schedules are created under Standard Time, the schedules continue to execute at the originally scheduled Standard Time, even though the storage modules are operating under Daylight Savings Time.

For example, if a schedule is configured under Standard Time to run at 2:00 PM, then the schedule initially runs (under Standard Time) at 2:00 PM. Then, when the local time changes to Daylight Savings Time, the schedule starts running at 3:00 PM instead of 2:00 PM. This is happening because the schedule is operating as if Daylight Savings Time doesn't exist; so the schedule continues to execute at 2:00 PM Standard Time. The Storage System Software does not include automatic adjustments for Daylight Savings Time.

Workaround

If you want snapshot schedules to operate at the same relative time all year, you must manually edit the schedules when the time changes in the spring and autumn.

4.9.2 When Running DSM for MPIO, Wait 60 Seconds before Re-logging Onto A Volume

Issue

If you are running Storage System Software DSM for MPIO and you log off a volume and immediately log back on the volume, you may find that the volume does not appear on the Disk Management tab of the Storage configuration category.

Workaround

After logging off a volume while running Storage System Software DSM for MPIO, wait for about 60 seconds before logging back on to the volume.

4.9.3 Auto Grow Not Converted to 6.6 Functionality When Upgrading

Issue

Auto grow has changed from version 6.5 to version 6.6. When you upgrade, volume configured with auto grow is set to Enabled in version 6.6. You cannot tell if it is using the manual algorithm or automatic auto grow.

Workaround

- Use the scripting command to determine the auto grow setting.

  Volume_autogrow_get <volume name> [<failure timeout seconds>]

  See the scripting chapter in the SAN user manual for more information about using scripting.

- You can reset auto grow on the volume if necessary. Simply use the Console to disable auto grow and reenable auto grow.

4.9.4 If iSCSI Load Balancing Is Misconfigured, After Upgrading To Version 6.6 Volume Operations Fail

Issue

After upgrading from version 6.5 to version 6.6, various volume operations fail. Snapshot schedules appear to stall, iSCSI log ins do not work, etc. The following error message displays:

"A Virtual IP is required. A VIP is required when linking or moving components to a load-balancing authentication group. Edit Cluster to enable the VIP."

Solution

Before upgrading, ensure that any clusters containing volumes in load-balanced authentication groups have a VIP configured. Clearing the load-balance flag on the authentication group(s) or configuration a VIP on the cluster(s) before the upgrade prevents the problem.

Workaround

However, you have already upgraded, there are a couple of solutions to try:

- Configure a VIP on the cluster that is causing the problem.
- If there is only one load-balanced authentication groups, clearing the load balance flag will fix the problem.

4.9.5 Volume Not Added to Volume List Appears in iSCSI Initiator

Issue

You create a cluster and configure the cluster to use iSNS. You then create a volume but do not add the volume to a volume list. The volume appears as a target in the iSCSI initiator. However, if you attempt to log on to this target, you receive an Authorization Failure message. This is a function of iSCSI discovery.

Solution

If you need to log on to the volume, add it to a volume list and create an authentication group, as described in the user documentation.

4.9.6 Volume Lists Must Contain Only Authentication Groups with Same Load Balancing Configurations

Issue

If a volume list contains one authentication group with load balancing and one authentication group without load balancing, it may not be possible for both of two cluster iSCSI clients to connect to the volume at the same time.

Solution

Only add authentication groups with the same load balancing configuration to a volume list.

4.9.7 Enable or Disable Load Balancing On An Authentication Group Requires Logging Off And Re-logging On To Volumes

Issue

The user changes the Enabled Load Balancing confugration of an authentication group. After-wards, some iSCSI clients may not ber able to reconnect to volumes because of the changes.

Workaround

Log off all iSCSI connections and log back on to reset the connections properly.

4.9.8 Sorting Volumes By Description Column In Ascending Order Corrupts Display

Issue

In the Console select Volumes in the navigation window. On the Volumes Details tab, click the “Description” column header for ascending sorting order. After a couple of seconds the following error is displayed: “String index out of range: -1” and the screen is unreadable.

Workaround

To correct this, once again select volumes and click another column header within a few seconds to go back to normal.

4.9.9 When Setting A Snapshot Schedule, Unable To Set Snapshot Time Between 12:00 A.M. And 12:59 A.M. (5836)

Issue

When setting a snapshot schedule in the Console to start between 12:00 a.m. and 12:59 a.m., the time changes to p.m. You cannot set the schedule to start during that first hour of the day.

Workaround

Choose another time to start the snapshot schedule.

4.9.10 After Taking Snapshots Of Multiple Volumes With Long Names, Storage Modules Appear To Return To Available Pool (5742)

Issue

The alert indicating the volume threshold change caused by taking snapshots, when triggered by multiple snapshots with long names causes storage modules to appear to return to the Available pool.

Workaround

Use the Find Modules wizard or the Find menu to search for storage modules. They will then return to the appropriate management group.

4.10 Remote IP Copy

4.10.1 Remote Copy Schedules Failed To Recreate When Restoring A Remote Management Group.

Issue

Two management groups, A and B, had a remote copy schedule copying from A to B. Both management group configurations had been backed up.

1. Management group A goes down and is restored from the configuration backup.
2. Later, management group B goes down and is restored from backup but now the remote schedule is lost.

Workaround

After restoring management group a, back up both management group configurations again. Otherwise you must manually re-enter the remote copy schedule.

4.10.2 After Restoration of Primary Management Group, Remote Snapshot Schedule on Remote Management Group Gets Deleted.

Issue

After performing a backup and restore of a primary management group with remote snapshots and remote snapshot schedules, the remote snapshot schedules on the remote management group are not properly restored. The schedule on the primary side still exists but is faulty, as if has not corresponding schedule on the remote side. Schedule remote copies will fail to start.

This failure to properly restore the schedule appears to happen in two situations.

- When there is a delayed start on the remote snapshot schedule. For example, when the remote snapshot schedule is setup, the initial start time is set at some point in the future, such as the next day
- When there are multiple remote snapshot schedules to restore.

Workaround

Delete the old remote snapshot schedules on the primary management group and recreate them. Scheduled remote snapshots will then resume.

4.11 ISCSI

4.11.1 Adaptec HBA Unable To See Target

Workaround

Do not use MS iSCSI initiator with the Adaptec HBA

4.11.2 iSCSI Closes All Shares After Reboot

Issue

If your iSCSI volumes are used by automatically-started Windows services (e.g., File Sharing), you must use the Microsoft* Initiator's “Bind Volumes” operation to make sure that those volumes are available before the services that require them are started.

Workaround

See the Microsoft* support article 870964 on the Microsoft support web site.

Also, see the section entitled “Running automatic start services on iSCSI disks” in the Microsoft* iSCSI Initiator Users Guide for more details.

4.11.3 An iSCSI Volume That Becomes Unavailable For Approximately 60 Seconds Or Longer May Cause Data Loss

The Windows Registry has a default maximum hold time setting of 60 seconds before a Microsoft Windows* system terminates a connection to an iSCSI device that is unavailable.

Therefore, an iSCSI volume that becomes unavailable for longer than 60 seconds may cause delayed write failures and potential data loss.

Solution

Change the Windows Registry setting for the default Maximum Request Hold Time to a very large (infinite) value.

Important: Back up your registry before making any changes.

1. Run regedit.exe.
2. Navigate to HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\4D36E97B-E325-11CE-BFC1-08002BE10318\####\Parameters (where #### is the index of the Microsoft* iSCSI initiator in the set of SCSI and RAID Controllers).
3. Double-click MaxRequestHoldTime in the right-hand pane. The Edit DMWord Value window opens.
4. Change the Base to decimal.
5. Enter a value of 600.
6. Click OK.
7. Save your changes and exit the Registry.
8. Reboot the system.

4.11.4 When Mounting Existing iSCSI Volumes On Different Servers, Volumes May Be Assigned Duplicate Drive Letters Or No Drive Letters

Issue

An iSCSI volume that was mounted on a server and assigned a drive letter is logged off from Server 1. It is then mounted on Server 2. Sometimes it picks up a drive letter that is already in use on Server 2. Sometimes it is not assigned a drive letter. The volume then becomes inaccessible.

Workaround

Open the Windows Disk Management console and assign a new drive letter to the volume. The volume should then appear in the directory structure.

4.11.5 Linux-iSCSI Initiator Cannot Reboot When Storage System Software Volume is Unavailable

The iSCSI Device Manager hangs when network problems prevent it from communicating with an SSR212MA. Because the default timeout for the Linux-iSCSI initiator is infinite, the initiator cannot reboot when it is unable to access the iSCSI volume on the SSR212MA.*Workaround*

Restore full network connectivity between iSCSI initiators and SSR212MAs. If this is not possible, disconnect the SSR212MA that the initiator can't communicate with from the network. Disconnecting will cause the managers to tell the client that it should stop attempting to contact that SSR212MA.

4.11.6 If Changing Permissions On An iSCSI Volume, Log On To A New Initiator Session To Complete The Changes

Issue

An iSCSI volume is mounted as a read/write volume and is in use.

You change the access permissions to read-only for the authentication group in the Console.

The permissions have not changed for the clients that are accessing the volume. They are still able to write to the volume.

Solution

To complete the process of changing permissions, you must log off the current initiator session for that volume and log on to a new session.

4.11.7 Microsoft* iSCSI Initiator Does Not Support Dynamic Disks

Issue

The Microsoft iSCSI initiator software does not support dynamic disks.

Workaround

Do not create dynamic disks to be used with the Microsoft iSCSI initiator.

4.11.8 iSCSI Volume Disappears From The iSCSI Initiator “Active Sessions” Window When Using Scheduled Snapshots

Issue

If you are using scheduled snapshots with an iSCSI volume, and the snapshot hard threshold is set to less than the volume hard threshold, the iSCSI volume disappears from the initiator Active Sessions window when the snapshot hard threshold is exceeded.

To recover from this situation:

1. In the Console, edit the snapshot schedule to increase the hard threshold.
2. Re-log in to the volume in the iSCSI initiator.

Workarounds

1. In the snapshot schedule, set the snapshot hard threshold to the same value as the volume thresholds, or
2. Use the auto_grow scripting feature to configure automatic threshold increases for the volume hard thresholds.

4.11.9 Unable to Build Oracle Application Cluster on iSCSI Raw Devices

Issue

If more than one iSCSI initiator attempts to access the same volume, the file system on the volume may become corrupted.

Workaround

Do not use volumes in a clustered iSCSI node configuration. Configure one iSCSI initiator per volume.

4.11.10 SSR212MA Failover Works As Long As A Virtual IP Address Is Used With iSCSI Initiators from Microsoft*, Intel®, Solaris*, Qlogic*, Adaptec*, Novell*, HP* and IBM*

Issue

To take advantage of the Storage System Software failover functionality in the iSCSI initiators from the listed companies, use a Virtual IP address when configuring clusters in the Console.

Workaround

Adaptec* Initiator

To ensure iSCSI volume availability in case of failover, it is recommended that the Session Recovery Timeout be set to 600 seconds. This is done using the Adaptec* iConfig utility.Qlogic*Initiator

To ensure iSCSI volume availability in case of failover, the following initiator configuration parameters must be set via the “Config Parameters” button on the Target Settings tab:

1. Default Timeout: 600 seconds
2. Connection Keep Alive Timeout: 600 seconds

Microsoft* Initiator

To ensure iSCSI volume availability in case of failover, the following registry key must be set:

```
[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-
11CE-BFC1-08002BE10318}\0000\Parameters]
```

MaxRequestHoldTime: 600 seconds

Note: The “0000” before “Parameters” in the registry path might vary. It could be 0001, 0002, etc. Search for MaxRequestHoldTime to find the key.

4.11.11 After Power Cycle, Load Balancing Does Not Distribute Requests Properly From A Microsoft Cluster

Issue

A Storage module is powered off and then powered on, and another storage module in the Storage System Software cluster handles all the connections to the volumes connected to that cluster. When the storage module is powered on again, load balancing does not redirect I/O to that storage module.

Workaround

1. Take one of the MS Cluster group offline.
2. Disconnect the iSCSI connection on both storage modules.
3. Reconnect the targets on both storage modules.
4. Bring the MS Cluster group back online.
5. Repeat steps 1-4 for all MS Cluster groups that host Storage System Software iSCSI disks.

4.11.12 When Using Storage System Software DSM for MPIO, If User Logs Off A Target Session While Client Is Accessing A Volume, The Path Disappears From MS Initiator

Issue

In the MS iSCSI initiator, when you open the Details window from the Targets tab and click on the Connections button on the Sessions tab, the number of iSCSI sessions displayed for a given volume is fewer than expected.

Workaround

1. Quiesce Client activity to the volume.
2. Completely log off the volume and reconnect.
3.

4.11.13 iSCSI Load Balancing Does Not Properly Balance iSCSI Sessions When Running A Mix of Servers with Storage System Software DSM for MPIO And Servers with iSCSI Load Balancing Enabled

Issue

A mixture of servers with Storage System Software with DSM for MPIO (Server Group-1) and servers with iSCSI Load Balancing enabled (Server Group-2) are accessing volumes in a storage cluster. iSCSI sessions from Server Group-2 are not properly load-balanced.

Workaround

If practical, the problem can be avoided by partitioning the Management Group into storage clusters such that a given cluster is accessed by only DSM hosts or only –DSM hosts.

4.11.14 iSCSI Load Balanced Connections Using Virtual IP Are Not Re-Assigned After Volume Is Migrated or Storage Module Removed from Cluster

Issue

After Migrating a volume or removing a storage module from a cluster, load-balanced iSCSI sessions are not re-assigned. While iSCSI connectivity is maintained throughout these operations, performance may not be optimal.

Workaround

1) Quiesce client activity to the volume.

2) Completely log off the volume and reconnect.

4.11.15 RedHat: Changing Authentication Type Causes Existing iSCSI Devices To Be Renamed

Issue

You configured an authentication group for iSCSI access. You then changed the access configuration, either to require CHAP or to remove or change CHAP requirements. After the change, the existing iSCSI devices are renamed and cannot be remounted.

Workaround

To change the authentication type of any volume (LVM or otherwise)

1. Unmount volumes and stop iSCSI services

```
# /etc/init.d/iscsi stop
```

2. Make appropriate changes to the authentication group (i.e. change from iqn to CHAP)
3. Make appropriate changes to the initiator (i.e. settings in /etc/iscsi.conf)
4. Start iSCSI services and remount volumes.

For LVM volume groups, the following steps are recommended since the system allows iSCSI services to be stopped even though iscsi_sfnet driver is still in use by the volume group.

To change authentication type of volumes being used in a volume group

1. Unmount volume/volume group

```
# umount /iSCSI
```

2. Deactivate the volume group

```
# vgchange –a n vgiSCSI
```

3. Stop iSCSI services

```
# /etc/init.d/iscsi stop
```

Then change to use CHAP or whatever authentication you want to test next. Then restart things in the reverse order:

```
# /etc/init.d/iscsi start
# vgchange –a y vgiSCSI
# mount/dev/vgiSCSI/lvol0/iSCSI
```

4.11.16 iSCSI Load Balanced Connections Using Virtual IP (VIP) Are Not Re-assigned After Volume Is Migrated or Storage Module Removed From Cluster

Issue

After migrating a volume or removing a storage module from a cluster, load-balanced iSCSI sessions are not re-assigned. While iSCSI connectivity is maintained throughout these operations, performance may not be optimal.

Workaround

1. Quiesce client activity to the volume.
2. Completely log off the volume and reconnect.

4.11.17 Failed Initiator Session Records Are Not Always Removed From Database

Issue

When a host disconnects ungracefully from a Storage System Software volume (e.g. the network fails, the host hardware is rebooted), the Console shows the iSCSI session for that host as Failed. If the host does not reestablish the connection within a day, the session is supposed to be considered dead and removed from the Console's iSCSI session displays. However, the failed sessions are usually not removed, resulting in lists of failed sessions displaying in the Console.

Workaround

None, However, as long as your volume is connected and accessible, and shows a connected session in the Console, you can ignore the failed session.

4.11.18 An Extra Microsoft iSCSI Session Is Created In The Console After Rebooting The Host

Issue

An extra iSCSI session is created in the Console after rebooting the host for the volume which is mounted with "Automatically restore this connection when the system boots" selected.

Explanation

This is a Microsoft issue in which different session IDs (iSCSI ISIDs) are used for the same host-volume pair, depending on how the session was established. After an ungraceful host shutdown, you might see duplicate iSCSI sessions in the Console, one with a Status of Failed and one a Status of Connected.

Workaround

Log off the automatically logged on persistent session and manually log back on to get rid of the spurious session.

4.11.19 Microsoft iSCSI Initiator Stops With Error

Explanation

In rare cases, the Microsoft iSCSI Initiator version 2.02 and 2.03 may stop after a storage module reboots.

Workaround

Manually restart the Microsoft iSCSI Initiator Service.

4.11.20 SuSE 9 and SuSE Linux iSCSI: Version 4.0.1-88.26 Initiator Reports Incorrect Driver State

Solution

Use the iSCSI initiator provided with the SLES 9 distribution.

4.11.21 Periodic iSCSI Event 39 Errors in Windows Host Event Log After Recoverable SAN Failure

Issue

Windows Event logs show long sequences of iScsiPrt event id 39 (task management commands sent) every 30 seconds accompanied by iScsiPrt event id 27 (no match for tag) and event id 9 (target did not respond in time) while IO is flowing to a volume. This could indicate that the MS iSCSI session is in the state where it issues spurious LUN resets.

Workaround

1) Quiesce the application that is using the volume
2) Log off the volume and log back on
3) Restart the application

Or

Reboot the application server

4.11.22 2-way CHAP Does Not Work With Solaris 10

Issue

Volume associated with an authentication group configured for 2-way CHAP cannot be mounted on Solaris 10.

Workaround

Use 1-way CHAP or no CHAP with Solaris !0.

4.12 Configuration Backup and Restore

4.12.1 SSR212MA Post-Install Qualification of Restored SSR212MA Stalls If Restored SSR212MA Has Different IP Address than That of Original SSR212MA

Issue

Back up an SSR212MA configuration file (SSR212MA-1). SSR212MA-1 becomes unavailable and you restore the backed up configuration of SSR212MA-1 to a second SSR212MA on the network (SSR212MA-2). SSR212MA-2 has a different IP address than the unavailable SSR212MA-1. As part of the post-install qualification, the Console searches for the newly configured SSR212MA-2 on the network. However, it is searching for the original IP address of SSR212MA-2 instead of the IP address that was saved in the SSR212MA-1 configuration back-up file. That search never completes because the IP address on SSR212MA-2 has changed and is now the IP address of SSR212MA-1.

Note: Restoring multiple SSR212MAs from a single backup file causes an IP address conflict.

Solution

Before restoring a backed-up SSR212MA configuration file, make certain that the new SSR212MA is configured with the IP address of the original SSR212MA.

Workaround

If the backed up configuration has been restored and the post-install qualification process can't complete because it cannot find the SSR212MA on the network, do the following:

1. On the Post install qualification window, click Cancel All Installs.
2. Search for the SSR212MA on the network using the correct IP address, or Find by Subnet and Mask.

4.12.2 If IP Address on SSR212MA Is Changed Using the Configuration Interface, Some Processes Continue to Use the Old IP Address

Issue

An SSR212MA in a management group has an IP address assigned. That IP address is changed using the Configuration Interface instead of using the Console. The new IP address is not universally updated in the Storage System Software and some functions continue to use the old IP address.

Workaround

To finish updating the IP address using the Console:

1. Log in to the SSR212MA with the new IP address.
2. In the SSR212MA Configuration Interface, navigate to the TCP/IP Network category.

On the Communication tab, click Update to synchronize the IP addresses of all managers.

4.12.3 Single Disk Errors Are Not Recovered In Clusters with SSMs Running Mixed SAN Software Versions

Issue

Release 6.3 contains functionality to recover from any single disk unrecoverable data error. This recovery functionality only works on SSR212MA's in clusters where all SSR212MA's are upgraded to version 6.3. If a cluster has one or more SSR212MA's running an earlier version of the software, than the recovery functionality will not work.

Workaround

Upgrade all SSM's to release 6.3 SAN software.

Fix

None.

4.13 MSCS

4.13.1 MSCS Cluster Failover While SSR212MA Cluster under Heavy Load Takes MSCS Cluster Off-line

Issue

If an MCS cluster failover occurs while the SSR212MA cluster is under very heavy load, the MCS cluster does not come back online until the load on the SSR212MA cluster decreases.

Workaround

Increase the "pending timeout" of each of the disk resources on the MCS cluster to the same as the "maxrequestholdtime" of 600.

Do the following on each "physical disk" resource that is actually an iSCSI disk on the SSR212MA.

1. Right-click on the disk in the MCS cluster administrator.
2. Select Properties > Advanced tab.
3. Change the "pending timeout: seconds" from 180 to whatever you used as a "maxrequestholdtime" for iSCSI in the registry.

4.13.2 Microsoft DSM Does Not Establish Expected Sessions In A MS Cluster Environment (5845)

Issue

The Microsoft DSM is not establishing all of the iSCSI sessions that are expected (1 for each storage module and 1 administrative session). Additionally, may experience problems failing the cluster resources back and forth between the Microsoft Nodes. This may result in loss of access to data and possibly data loss.

Workaround

1. Log off all iSCSI connections.
2. Un-install the Microsoft DSM.
3. Remove all persistent sessions.
4. Reboot the host.
5. Run the Scrubber utility, available from Microsoft and described at http://support.microsoft.com/kb/277222.

4.14 Dell OpenManage Secure Port Server

4.14.1 Unable To Install or Load Console with Dell's Secure Port Server Service Started

Issue

Using Microsoft Windows* on a Dell* Server with the Dell* OpenManage Secure Port Server service, the user cannot properly install the Console or start the Console.

Workaround

Stop the Dell* OpenManage Secure Port Server service when installing or running the Console.